

## **IBM DB2 12 for z/OS Tools and Utilities**

*New Feature Support and the Impact of Continuous Delivery*

By Craig S. Mullins

President and Principal Consultant at Mullins Consulting, Inc.

## Executive Summary

IBM® DB2® 12 for z/OS® delivers a lot of new functionality. And it is vitally important that your DB2 tools and utilities are ready to work with and manage your databases and applications that use this new functionality. Managing DB2 without the requisite management tools and utilities is error-prone; it raises the risk of failure or data integrity issues. Therefore, if the DB2 tools and utilities you are using do not support new features, it can delay the use of new DB2 functionality, migration to the new version of DB2, or destabilize DB2 environments.

Not upgrading to a new version of DB2 because your tools vendor does not adequately support new version functionality is also a risk because it delays implementation of new features that improve performance and minimize development effort. For these reasons, it is wise to look for a vendor that can quickly support new versions of DB2, so that you can migrate to a new version of DB2 and all the supporting tools and utilities exploit new features on your desired schedule.

IBM continues to provide prompt support for new DB2 functionality to enable all DB2 customers to immediately benefit from new version features and functionality.

This white paper examines IBM's support and exploitation of DB2 12 for z/OS in its tools and utilities showing a high level of support and preparation that can ease your transition to the newest version of DB2 for z/OS.

## DB2 12 for z/OS

DB2 12 for z/OS was released for general availability on October 21<sup>st</sup> 2016 and organizations are looking to migrate to the latest and greatest version of DB2 to take advantage of the many new features and enhancements it provides. Of course, organizations will migrate to a new version based on many factors, including their tolerance for change, desire for specific features and the support available for the new version in each organization's applications, tools, and utilities.

Tool and utility support is particularly important and is the focus of this white paper. Modern enterprise applications rely on DB2 for z/OS to store and maintain the data that is the lifeblood of today's organizations. But the core DB2 software does not deliver all of the functionality needed to support, manage, and administer large-scale database development. Management tools and utilities that enhance the functionality of DB2 are required to deliver robust performance, round-the-clock availability, and scalable applications. For this reason, it is important that the tools and utilities used to manage DB2 support the features and functionalities of a new version before your organization fully migrates to the new version.

As such, organizations will be looking for DB2 tools that support the DB2 12 for z/OS features that they will be adopting and relying upon. There are hundreds of new features and enhancements in the new version. These improvements are the reasons that organizations want to migrate to the new version, but there are some features that top the list of things that DB2 users will adopt first. Therefore, this white paper will examine the top six DB2 12 for z/OS features that will immediately impact most organizations and examine how IBM's DB2 tools and utilities exploit and manage these new features.

I chose six significant features that DB2 customers will be experiencing or looking to exploit first. I based the selection of these features on my consulting practice and experience dealing with DB2 implementations of many different types and sizes. The features we will examine are:

- Partition-by-growth Relative Page Numbering Universal Table Spaces (PBG RPN UTS)
- Index Fast Traverse Blocks (FTBs)
- Temporal Improvements
- Dynamic/Static SQL Management
- Online Schema Enhancements
- Continuous delivery

The good news is that IBM has a consistent track record of promptly supporting new functionality in its DB2 tools and utilities. Customers rely on IBM's utilities to manage and move data, and as such IBM's utilities must be ready to support all new DB2 version functionality on day one. The tools to manage DB2 are just as important as the utilities, and therefore IBM works to support new features quickly. For example, the IBM performance team relied upon **OMEGAMON® XE for DB2® Performance Expert** on z/OS® (OMPE) during the development cycle. As such, the product had to be up-to-date to enable the team to do its job.

As you read through this document it will become clear that IBM's DB2 tools and utilities are ready to support your DB2 12 environment today. But to be more forward-looking, you will need to re-evaluate your typical rollout policies and procedures for new DB2 versions because of IBM's move to Continuous Delivery for new DB2 functionality. Changes to DB2 will become available in the future without a new version of the software and, as such, it becomes even more important that your DB2 tools and utilities will be there with support for that functionality quickly. Acquiring tools from the same vendor that makes the DBMS

can help you to minimize the confusion and downtime that can result from trying to use management tools and utilities that are incompatible with a current DB2 level.

This will become more clear as you read through the subsequent sections of this white paper.

## **Partition-By-Range Relative Page Number Universal Table Spaces** (or PBR RPN UTS)

### ***What is it?***

A significant new feature for supporting big data in a DB2 12 environment is relative page numbering (or RPN) for range-partitioned table spaces. An RPN range-partitioned table space can be created, or an existing range-partitioned table space can be changed to RPN via an ALTER TABLESPACE with PAGENUM RELATIVE, followed by an online REORG of the entire table space.

An RPN table space delivers many benefits for availability and storing large amounts of data. This requires an expanded RID, which increases from 5 bytes to 7 bytes.

From an availability perspective, you can specify DSSIZE at the partition level for RPN table spaces. Furthermore, the allowable DSSIZE value is no longer dependent on the page size and number of table space partitions. The DSSIZE change can be an immediate change (no online REORG required to take effect) as long as the change does not decrease the DSSIZE value. You still can decrease DSSIZE, but only at the table space level.

From a big data perspective, the DSSIZE can grow up to 1 TB for a partition. And the maximum table size increases to 4 PB with approximately 280 trillion rows per table. That is a lot of data that can be stored. Think about it this way: if you were to insert 1000 rows per second it would take more than 8000 years to fill the table to capacity!

### ***What does it mean for IBM tools and utilities?***

A new type of table space format has a big impact on your DB2 tools and utilities. First of all, utilities that read and manipulate data need to understand the new page format and layout. All of the IBM DB2 utilities are fully compatible with PBR RPN table spaces. Remember, too, that RIDs expand to 7 bytes from 5 for RPN table spaces. Therefore, a REORG of a PBR PRN table space with SHRLEVEL change requires a mapping table with two larger columns: SOURCE\_RID CHAR(7) (formerly 5 bytes) and TARGET\_XRID CHAR(11) (formerly 9 bytes). The IBM REORG utility fully supports this.

The CHECK DATA utility that uses an exception table with a RID column must have a 7-byte RID column (formerly 5-byte column). Since 7 bytes is acceptable for all objects, on a CHECK DATA utility where "Include RID column" =Y has been specified in the DB2 Automation Tool, it will always generate a mapping table with a 7-byte RID column.

But utilities are not the only type of DB2 tools that need to be adapted to support PBR RPN table spaces. The IBM DB2 Automation Tool also fully supports PBR RPN table spaces. This includes using mapping tables with the larger columns (which are usable by all types of table spaces), as well as understanding all the new keywords and parameters. An example is with the RECOVER utility where PAGE is specified and DSNUM must also be specified for table spaces or indexes with relative page numbers.

Another tool that has been augmented to supported PBR RPN table spaces is the IBM DB2 Cloning Tool, which enables fast, efficient cloning of either subsystem or application data. It fully supports the larger PBR RPN table spaces, and can be used to clone them. It also delivers additional DB2 12 improvements for new DDL attributes, support for multiple copy pools and operating on compressed LOB data.

The DB2 Admin Tool also provides support for creating and altering PBR RPN table spaces, as shown in the screen capture in Figure 1. The MAKEPBR2 option enables this capability.

```

DB2 Admin ----- DB1D Redefine Table Space ----- Row 1 to 1 of 1
Command ==> █                                         Scroll ==> PAGE

Commands: NEXT ORIGINAL                               MAKEPBG MAKEPBR MAKEPBR2
Line commands: S - Split part   R - Remove part 0 - Original data
                C - Clear data
CREATE TABLESPACE: DSN8S12Q IN DSN8D12P

Owner . . . . . > Owner type . . (U/R)

Numparts . . . . . 0                               LOB . . . . . NO
Define . . . . . YES                               LOG . . . . . YES
Member Cluster . . NO                               SEGSIZE . . . . 4   CCSID . . . . EBCDIC
Buffer Pool . . . . BP0                             Close Rule . . NO   Max Rows . . 255
Lock Size . . . . . PAGE                           Lock Part . . . NO  Lock Max . . 0
Max Partitions . . 0                               PAGENUM . . . . .

S   Part      Pqty      Sqty  Free   E T S
   Page PF  Cmp R M T VCAT      Stogroup GBPCach DSSIZE
----->-----
      0         -1       -1    0  5 NO  N Y I DSNDCAT  DSN8G120 CHANGED
***** END OF DB2 DATA *****

```

**Figure 1.** The IBM DB2 Admin Tool with the MAKEPBR2 option.

IBM DB2 Recovery Expert, which uses automation to reduce recovery think time and execution time, also fully supports PBR RPN table spaces. And, as previously mentioned, the IBM DB2 utilities driven by Recovery Expert (COPY, RECOVER and REBUILD) all fully operate on PBR RPN structures. IBM DB2 Recovery Expert supports

dropped object recovery for all new database object attributes including PBR RPN parameters.

From a performance perspective, OMPE supports PBR RPN table spaces in the reports and displays that show locks owned, locks waited on and lock conflicts.

## **Index FTBs**

### ***What is it?***

Many new in-memory processing capabilities have been added in DB2 12 for z/OS. Perhaps the most interesting are the new Fast Traversal Blocks, or FTBs, which are in-memory structures that can be used with unique indexes. DB2 detects which indexes are frequently used for traversals, and when a threshold is hit DB2 will build an FTB for the index in a storage area outside the buffer pool. This causes the top levels of the index to be cached thereby making it efficient to perform very fast traversals through the cached levels of the index.

FTBs are either on or off for the entire subsystem and this is controlled using the INDEX\_MEMORY\_CONTROL DSNZPARM. The AUTO setting, which is the default, indicates that 500 MB or 20 percent of the buffer pool will be used for FTBs (whichever is larger). Alternatively, you can set the upper limit to a number between 10MB and 200 GB, or you can DISABLE the feature altogether.

This feature is likely to be used by customers having applications that perform frequent lookups where the index is used predominantly for reads. In those scenarios FTBs may be able to deliver a significant performance improvement.

### ***What does it mean for IBM tools and utilities?***

There are two new DB2 IFCID trace records that report on index FTB usage in DB2 12 for z/OS: IFCID 389 and 477. IFCID 389 traces indexes with FTB structures and IFCID 477 traces allocation and deallocation of FTB structures. OMPE provides the details of this FTB trace information in the RECTRACE report. An example of the information provided by OMPE for IFCID 389 is shown in example in Figure 2.

CONNECT CORRNAME CORRNMBR	INSTANCE CONNTYPE	END_USER RECORD TIME TCB CPU TIME	WS_NAME DESTNO ACE	IFC ID	DESCRIPTION	TRANSACTION DATA
N/P	D0B9113C791A	N/P	N/P			N/P
N/P	'BLANK'	18:19:12.26938222	12133	1 389	FTB INDEXES	NETWORKID: DSND20G LUNAME: DSND20G
N/P		N/P				
NUMBER OF INDEXES :		7				
DBID: 265	PSID: 11	PARTITION NO:	1	INDEX LEVELS:	2	FTB SIZE: 2339584
DBID: 265	PSID: 18	PARTITION NO:	1	INDEX LEVELS:	3	FTB SIZE: 1368064
DBID: 265	PSID: 18	PARTITION NO:	2	INDEX LEVELS:	3	FTB SIZE: 1560064
DBID: 265	PSID: 18	PARTITION NO:	4	INDEX LEVELS:	3	FTB SIZE: 5209600
DBID: 265	PSID: 18	PARTITION NO:	5	INDEX LEVELS:	2	FTB SIZE: 2560000
DBID: 265	PSID: 16	PARTITION NO:	2	INDEX LEVELS:	2	FTB SIZE: 146176
DBID: 265	PSID: 16	PARTITION NO:	3	INDEX LEVELS:	3	FTB SIZE: 195584

**Figure 2.** Index FTB trace information.

As you can see, this information shows the number of indexes with FTBs along with the number of levels in the FTB and the size of the structure. Such details will be important for DBAs looking to manage and support index FTBs in DB2 12.

Additionally, the OMPE batch SYSPARMS report shows the DSNZPARM settings for INDEX\_MEMORY\_CONTROL, which shows the amount of memory that DB2 should allocate for fast traversing of unique indexes. Utilities such as IBM REORG and REBUILD INDEX support FTB processing. Additionally, the IBM DB2 Admin Tool helps you manage ZPARMS and can also be used to control FTB processing.

## Temporal Improvements

### *What is it?*

There are several new features in DB2 12 for z/OS that simplify the usage and management of temporal data in DB2 including temporal referential constraints, more flexible time periods, logical transactions, and RTS history in the DB2 system catalog.

The most significant of these is probably the addition of *temporal referential constraints*. Prior to DB2 12, temporal RI required triggers or stored procedures to implement. But with DB2 12, you can define a temporal referential constraint for a temporal table that contains a BUSINESS\_TIME period by specifying the PERIOD BUSINESS\_TIME clause in the definition of the constraint.

The second new temporal feature is improved *flexibility for defining the application period* in temporal tables. The application period is defined with two date/time columns, one specifying the beginning of the period and the other specifying the end of the period. Prior to DB2 12, the beginning value of a period had to be inclusive, and the end value had to be exclusive. So for a period beginning at 2017-01-01 and ending 2019-10-01, 2017-01-01 is part of the period but 2019-10-01 is not part

of the period. In DB2 12, you can create an application-period temporal table with a BUSINESS\_TIME period that is inclusive-inclusive. That means that the end value is considered to be part of the period, instead of outside the period definition.

DB2 12 also offers *temporal logical transactions* on system temporal tables. With this new feature, DB2 supports logical units of temporal work that are not determined by COMMIT and ROLLBACK; values for row-begin and row-end columns are determined by applications based on a built-in global variable that you can set.

The final significant temporal improvement in DB2 12 is the ability to maintain *RTS history* using DB2's system-time temporal capabilities. When history is enabled, both SYSIBM.SYSINDEXSPACESTATS and SYSIBM.SYSTABLESPACESTATS will have an associated SYSTEM\_TIME history table that will house historical RTS as new RTS data is accumulated.

**What does it mean for IBM tools and utilities?**

The impact of temporal improvements on DB2 tools largely involve DDL options and parameters, and therefore the impact centers on the IBM DB2 Admin Tool. The Admin Tool supports DDL changes for the DB2 12 temporal enhancements. Figure 3 shows an example of using the IBM DB2 Admin Tool using business time temporal tables with an inclusive-inclusive application period.

```

ADBP7TOP ----- DB2A ALTER - Table Options -----
Command ==>

New schema . : SYSADM
New name . . : TBADBE01

Enter table options below:

AUDIT . . . . . NONE          (None, Changes, or All)
DATA CAPTURE . . . . . NONE    (None/Changes)
VALIDPROC . . . . . NULL      (NULL/Program name)
EDITPROC . . . . .
  WITH ROW ATTRIBUTES . .      (Yes/No)
RESTRICT ON DROP . . . . . NO   (Yes/No)
VOLATILE . . . . . NO         (Yes/No)
APPEND . . . . . NO          (Yes/No)
LABEL . . . . .
COMMENT . . . . . >
Business period . . . . . YES   (Yes/No)
  Begin column . . . . . BUS_START > (? to lookup)
  End column . . . . . BUS_END  > (? to lookup)
  INCLUSIVE . . . . . YES      (Yes/No)
System period . . . . . YES    (Yes/No)
Versioning . . . . . YES      (Yes, No, or Chg)
ENABLE ARCHIVE . . . . . NO    (Yes, No, or Chg)
  
```

**Figure 3.** *Inclusive-Inclusive support in the DB2 Admin Tool.*

Furthermore, the IBM® DB2® Administration Tool for z/OS® (DB2 Admin Tool) supports RTS history tables with catalog visibility to the information they contain. This enables DBAs to use the DB2 Admin Tool to track statistical trends and patterns for their database objects and be better prepared to manage systems with rapidly changing data requirements.

## Dynamic/Static SQL Management

### *What is it?*

Most modern application development is done using dynamic SQL. But some features work only with static SQL and others only with dynamic SQL. DB2 12 for z/OS delivers functionality that minimizes the difference between static and dynamic SQL.

***Dynamic plan stability*** brings the plan stability feature of static SQL to dynamic SQL. With plan stability for static SQL, you can use the PLANMGMT parameter of REBIND to save off old access paths that can be switched back to active if the new access paths are inefficient for any reason.

DB2 12 introduces dynamic plan stability, which is a little different but should prove to be quite useful. You can choose to stabilize dynamic SQL access paths by storing them in the DB2 system catalog. There are options for selecting which queries to stabilize, so you do not have to store all dynamic SQL. This is controlled by the -START DYNQUERYCAPTURE command.

When a dynamic statement is run, DB2 will look in the dynamic statement cache first. If it is not found there, DB2 will look in the system catalog before resorting to a full prepare.

This option can make your dynamic SQL more static. But there is another new DB2 12 capability that can make your static SQL more dynamic: ***static Resource Limit Facility*** (RLF). The RLF, heretofore, could only be used to govern dynamic SQL statements. RLF tables, manipulated by DBAs, contain limits that make sure that dynamic SQL statements do not consume too many resources such as CPU, I/O, locks, etc. This enables improperly tested SQL or poor resource planning from disrupting performance.

But dynamic SQL is not the only type of SQL that could cause performance issues; static SQL transactions can benefit from the same controls. DB2 12 extends the RLF to support static SQL statements thereby improving the ability to avoid problem applications from dominating your system's resource consumption.

### *What does it mean for IBM tools and utilities?*

For dynamic plan stability, it is important to be able to track which plans have been stabilized in order to analyze access paths and dynamic SQL performance. OMPE has reports and online displays of the dynamic SQL cache that show if a plan has been stabilized.

Additionally, OMPE reports on dynamic plan stability issues that help to tune dynamic performance. OMPE details how dynamic plan stability is performing by showing a cache+catalog hit ratio and the number of short prepares with regard, to caching or loading from the catalog.

## **Online Schema Enhancements**

### *What is it?*

The ability to make changes to database structures without causing major outages is a key requirement of modern database applications. With its online schema change initiative, begun in DB2 V8, IBM has been steadily expanding the number and type of changes that can be made without an outage. And DB2 12 for z/OS adds even more online schema change capabilities. This includes:

- It becomes possible to recover to a point in time before materialization of pending definition changes for PBG table space. In past releases this was not possible.
- The ability to defer column-level ALTER operations such as changing data type, length, precision or scale which are, by default, immediate changes. This is done at the subsystem level by setting the DDL\_MATERIALIZATION DSNZPARM value to ALWAYS\_PENDING. By making all of these changes pending, it makes it easier to manage multiple changes without having to determine which are immediate and which are pending.
- When you issue ALTER INDEX to change compression for an index in a universal table space the change is now a pending change that places the index in AREOR status. The indexes can continue to be used and the new compression specification is materialized by a subsequent online REORG INDEX or online REORG TABLESPACE at the table space level.
- You can insert a new partition in the middle of a partitioning range and convert an existing PBR table space to partition by range with RPN.

- The TRANSFER OWNERSHIP command makes it possible to transfer the ownership of an object (to another user or a role) without having to DROP and re-CREATE the object.

### *What does it mean for IBM tools and utilities?*

Many of the scalability and availability benefits in DB2 12, which take advantage of greater scalability and application availability, use the IBM DB2 12 REORG utility. These enhancements only take effect after the next execution of the DB2 REORG utility, which means that the IBM REORG utility must support the new online schema enhancements in order for them to work in DB2 12. As such, you can rely on the IBM REORG utility to support online schema enhancements whenever such new functionality is added to DB2.

In DB2 12, you can also reorganize a single partition-by-growth (PBG) table space and have it automatically overflow to a new partition if the data no longer fits after the REORG operation.

And very importantly, the IBM DB2 Admin Tool fully understands the new online schema enhancements so that it can implement database changes in the most efficient manner – building plans that make use of immediate and online changes where supported by DB2 12. The current situation is that many different types of DB2 structures and parameters can be altered, and in multiple different ways, and in different orders that can impact how changes are implemented by DB2. As such, it is becoming vital for organizations to rely on a tool like the Admin Tool which can mask the complexity and implement any type of change using the best available approach.

The TRANSFER OWNERSHIP capability is supported on the catalog navigation panels for storage groups, databases, table spaces, tables, indexes, and views, so you can easily use Admin Tool to transfer the ownership of any database objects to another user.

The IBM DB2 Automation tool supports altering index compression as a pending definition change by adding the ability to check for the new pending definition changes via a new exception condition.

IBM DB2 Recovery Expert provides recovery through online schema changes to support point in time recovery before materialization of pending definition changes for PBG table spaces. In fact, Recovery Expert allows for recovery to a point-in-time prior to schema materializations that the native RECOVER utility does not support.

If the appropriate IFCID traces are started, OMPE has been enhanced so that it will show details of TRANSFER OWNERSHIP executions on the DB2 Record Trace report as well on the Thread Detail display. And, once again, OMPE provides support for the new

DDL\_MATERIALIZATION DSNZPARM parameter settings on the batch SYSPARMS report.

## Continuous Delivery

### *What is it?*

Perhaps the most significant change with DB2 12 for z/OS is IBM's movement to a continuous delivery model for future enhancements. This changes the long-standing model of new versions every 3 years, to a regular, step approach to delivering new functions that involves shipping changes and service (both preventative service and defect fixes) in the same service stream. Instead of a major disruption every 3 years or so, DB2 customers will be able to adopt new functionality in a controlled manner and at a flexible pace and time.

So instead of the various migration modes involved in past DB2 migration projects, DB2 12 introduces a function level that consists of a 9-byte string showing the version, release and modification indicator. For example, V12R1M500 indicates version 12 release 1 modification 500, which is the first function level that supports new DB2 12 functionality.

A new command, `ACTIVATE`, is used to activate a specific function level, after which all functionality supported at that function level and previous function levels become available to be used.

Organizations need to prepare for how they will roll out new function levels, including communication of changes, training staff on new features, testing, and so on. Furthermore, as new functionality is introduced through activation of new function levels over the life of DB2 12, you will need to understand how your tools and utilities vendors supports each new function level and coordinate your activities accordingly.

### *What does it mean for IBM tools and utilities?*

Changes will be introduced and made available in DB2 more rapidly than before, so it becomes even more important that your utilities and tools vendor is able to keep up and support these new features on a timely basis. If the tools you use to manage DB2 are not ready to support new features then the functionality is, for all intents and purposes, not available to you. If you cannot monitor, manage, administer, and track databases using new the features, then you increase the risk of unidentified performance issues and failures that cannot be tracked.

***Quick support of new DB2 functionality:*** The expectation is that organizations will apply maintenance faster than they will activate new function. In other words, the function level will lag behind the maintenance level. Nevertheless, the goal must be for

DB2 tools and utilities to rapidly integrate and support new functionality quickly, as it is introduced to DB2 for z/OS. This is even more critical in terms of avoiding regression of existing capability that customers rely upon today. It is IBM's intent to quickly provide support for new functionality in this new world of continuous delivery. Given its previous track record for support, it is reasonable to rely on IBM's tools and utilities for DB2 being ready with new function support when you are ready to activate that functionality.

IBM DB2 tools and utilities are key to helping you exploit new capability in DB2. That's why IBM makes both the DB2 utilities and tools available to customers in the Early Support Program (ESP). With DB2 12 for z/OS, real customers tested new functionality in DB2 during the ESP using IBM tools and utilities.

The new COMPRESSRATIO column in DB2 system catalog offers another good example of IBM DB2 utility support of Version 12. The purpose of this new column is to improve the accuracy for calculating sort work data sets. COMPRESSRATIO gets the average percentage of bytes saved when COMPRESS YES is specified. The value is based on an average row length and varies depending on the actual length of the data rows. The REORG, LOAD REPLACE and RUNSTATS utilities automatically update this new column.

This white paper should convince you that IBM is committed to keeping its tools and utilities up-to-date with new version support. And IBM is inherently suited to support DB2's future continuous delivery model. Here are just a few more of the immediate new version capabilities supported by IBM.

The DB2 Automation Tool supports all new utility enhancements (including all new parameters and keywords), index compression deferred alter, and RTS improvements. The new command ACTIVATE NEW FUNCTION is supported in the DB2 Command Processor of the Automation Tool. Additionally, the Automation Tool can generate REPAIR Catalog and REPAIR Catalog TEST for index spaces, which is a new Version 12 capability.

DB2 Recovery Expert delivers support for additional DB2 12 enhancements including support to skip recovery of unchanged objects, support for multiple copy pools and new recovery options, and dropped object recovery of all new object attributes.

OMPE offers wide support for DB2 12 performance features including Dynamic Plan Stability, buffer pool and EDM pool enhancements, raised partition limits, LOB compression, fast unclustered Insert, authorization enhancements, new IFCIDs, and new system parameters.

IBM Query Monitor for DB2 has been updated to support all DB2 12 control block changes. Many infrastructure changes were implemented to detect and properly handle

DB2 12 continuous delivery changes. And performance improvements were achieved by exploitation of new opcodes based on DB2 12 hardware requirements.

Indeed, every IBM tool and utility for DB2 has been updated and improved to support and where possible exploit, DB2 12 for z/OS.

***Continuous tool and utility functionality delivery:*** But there is another nuance to continuous delivery for tools and utilities. We see continuous delivery of many new features that add capabilities to the tools and utilities that are not driven by DB2 engine changes.

For example, IBM has delivered additional offload to **zIIP processors** to reduce the cost of running IBM DB2 utilities. As of DB2 12, the RELOAD phase of the LOAD and REORG utilities becomes zIIP eligible. This comes as yet another improvement in a long succession of zIIP offloading that started with DB2 V8:

- V8 Index build in LOAD, REORG and REBUILD
- V9 REORG UNLOAD
- V10 RUNSTATS
- V11 Inline stats
- V12 LOAD and REORG RELOAD

A good example of a tool with significant enhancements is the **DB2 Recovery Expert**. The tool recently added functionality for log-based object recovery, enabling DBAs to complete end-to-end recovery of a dropped object from the log. This includes recovery of DDL schema, data, DCL (authorizations), and rebinding of application programs. DB2 Recovery Expert also added support for redirected recovery, making it possible to restore to another system or different set of objects on same system. But that is not all; it also now delivers transaction recovery support so DBAs can generate and execute UNDO SQL for a single transaction or a range of transactions over time.

IBM additionally has made a significant investment in support for the **IBM® DB2® Analytics Accelerator for z/OS®** in its tools and utilities. There are many areas where tool and utility support for the Accelerator can benefit DBAs, but even more so when Accelerator-only tables (AOTs) are utilized because the data for these tables only exists on the Accelerator.

The use of the **IBM DB2 Analytics Accelerator Loader** with the Accelerator significantly expands customer's analytic capabilities with support for additional disparate non-DB2 data sources. As such, many DB2 customers have built large repositories of data in the Accelerator that cannot be easily recaptured in the event of failure. This can be operational data such as SMF or other types of log data, or application data originating in flat files, VSAM, IMS or any other data source. The data might be maintained and backed up in its originating system, but there is a significant



- Connect to DB2 for z/OS and navigate database objects, viewing details and links to related objects and object dependencies
- Browse and edit data
- Create, alter and drop database objects, such as of tables, indexes, constraints and table spaces
- Show system privileges from the perspective of group or user, role, or database objects

DSM makes it easier for all types of users to get the most out of their DB2 experience.

### **The Bottom Line**

Not every change in every IBM DB2 tool and utility is outlined in this document. The intent is not so much to detail in depth every change, but instead to show the types of changes that IBM has made and thereby show the focus and effort being put into maintaining the currency of the IBM DB2 tools and utilities.

After reviewing the details of how IBM supports these six Version 12 improvements, it should be obvious that IBM has put a lot of time and effort into expanding the functionality of their DB2 tools and utilities; both to support Version 12 and to deliver on-going new capabilities.

A good overview of all the new features in DB2 12 can be found in the What's New guide (link and number) and the IBM DB2 12 for z/OS Technical Overview redbook (SG24-8383 <http://www.redbooks.ibm.com/redbooks/pdfs/sg248383.pdf>).

You should examine each new feature that you are looking to implement to see how your tools and utilities provide support for your desired functionality. If your tools and utilities are not offering the support you need, when you need it, be sure to take a look at IBM's offerings. Because frequently their tools will be more up-to-date.



## About the Author

Craig S. Mullins is a data management strategist, researcher, and consultant. He is president and principal consultant of Mullins Consulting, Inc.

<http://www.mullinsconsulting.com>

Craig has been named by IBM as a Gold Consultant and an IBM Champion for Analytics. He was also named one of the Top 200 Thought Leaders in Big Data & Analytics by AnalyticsWeek magazine (<http://analyticsweek.com/top-200-thought-leaders-in-bigdata-analytics/>).

Craig has over three decades of experience in all facets of database systems development and has worked with DB2 since V1. You may know Craig from his popular books:

- *DB2 Developer's Guide, 6<sup>th</sup> edition* – containing more than 1500 pages of in-depth technical information on DB2 for z/OS
- *Database Administration: The Complete Guide to DBA Practices and Procedures, 2<sup>nd</sup> edition* – the industry's only comprehensive guide to heterogeneous database administration.

© Copyright Mullins Consulting, Inc. 2017

### **For more information**

To learn more about IBM DB2 utilities and tools, contact your IBM representative, or visit: [ibm.com/analytics/us/en/technology/db2/zos/tools/](http://ibm.com/analytics/us/en/technology/db2/zos/tools/)

### **Disclaimer**

IBM, the IBM logo, ibm.com, DB2, and z/OS are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.